Search Results

Search Results for: [tree <near> search <near> engine<AND>((direct <near> table<AND>((full <near> match<AND>((longest <near> match))))))]
Found 8 of 96,434 searched. → Rerun within the Portal

Search within Results

1 Scalable high-speed prefix matching 54%
Marcel Waldvogel , George Varghese , Jon Turner , Bernhard Plattner
ACM Transactions on Computer Systems (TOCS) November 2001
Volume 19 Issue 4

Finding the longest matching prefix from a database of keywords is an old problem with a number of applications, ranging from dictionary searches to advanced memory management to computational geometry. But perhaps today's most frequent best matching prefix lookups occur in the Internet, when forwarding packets from router to router. Internet traffic volume and link speeds are rapidly increasing; at the same time, a growing user population is increasing the size of routing tables against which p ...

2 Small forwarding tables for fast routing lookups 40%
Mikael Degermark , Andrej Brodnik , Svante Carlsson , Stephen Pink
ACM SIGCOMM Computer Communication Review , Proceedings of the ACM SIGCOMM '97 conference on Applications, technologies, architectures, and protocols for computer communication October 1997

Volume 27 Issue 4

3 Inverted files versus signature files for text indexing 12%
Justin Zobel , Alistair Moffat , Kotagiri Ramamohanarao
ACM Transactions on Database Systems (TODS) December 1998
Volume 23 Issue 4

Two well-known indexing methods are inverted files and signature files. We have undertaken a detailed comparison of these two approaches in the context of text indexing, paying particular attention to query evaluation speed and space requirements. We have examined their relative performance using both experimentation and a refined approach to modeling of signature files, and demonstrate that inverted files are distinctly superior to signature files. Not only can inverted files be used to ev ...

4 Scalable high speed IP routing lookups 7%

Marcel Waldvogel , George Varghese , Jon Turner , Bernhard Plattner
ACM SIGCOMM Computer Communication Review , Proceedings of the ACM SIGCOMM '97 conference on Applications, technologies, architectures, and protocols for computer communication October 1997

Volume 27 Issue 4

5 Field programmable port extender (FPX) for distributed routing and queuing 3%

John W. Lockwood , Jon S. Turner , David E. Taylor
Proceedings of the ACM/SIGDA international symposium on Field programmable gate arrays February 2000

Field Programmable Gate Arrays (FPGAs) are being used to provide fast Internet Protocol (IP) packet routing and advanced queuing in a highly scalable network switch. A new module, called the Field-programmable Port Extender (FPX), is being built to augment the Washington University Gigabit Switch (WUGS) with reprogrammable logic.FPX modules reside at the edge of the WUGS switching fabric. Physically, the module is inserted between an optical line card and the WUGS gigabit switch ...

6 Practical dictionary management for hardware data compression 1%

Suzanne Bunton , Gaetano Borriello
Communications of the ACM January 1992
Volume 35 Issue 1

7 External memory algorithms and data structures 1%

Jeffrey Scott Vitter
ACM Computing Surveys (CSUR) June 2001
Volume 33 Issue 2

Data sets in large applications are often too massive to fit completely inside the computers internal memory. The resulting input/output communication (or I/O) between fast internal memory and slower external memory (such as disks) can be a major performance bottleneck. In this article we survey the state of the

art in the design and analysis of external memory (or EM) algorithms and data structures, where the goal is to exploit locality in order to reduce the I/O costs. We consider a varie ...

8 NiagaraCQ 0%

Jianjun Chen , David J. DeWitt , Feng Tian , Yuan Wang
ACM SIGMOD Record , Proceedings of the 2000 ACM SIGMOD international conference on Management of data May 2000
Volume 29 Issue 2

Continuous queries are persistent queries that allow users to receive new results when they become available. While continuous query systems can transform a passive web into an active environment, they need to be able to support millions of queries due to the scale of the Internet. No existing systems have achieved this level of scalability. NiagaraCQ addresses this problem by grouping continuous queries based on the observation that many web queries share similar structures. Grouped querie ...

Results 1 - 8 of 8 **short listing**
